

EV355228665
IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Linguistic Disambiguation System and Method Using
String-Based Pattern Training To Learn To Resolve
Ambiguity Sites**

Inventor(s):
Eric Brill

ATTORNEY'S DOCKET NO. MS1-524USC1

TECHNICAL FIELD

This invention relates to linguistic ambiguity resolution. More particularly, this invention relates to systems and method for training linguistic disambiguators using string-based patterns.

BACKGROUND

Natural language processing systems are computer implemented software systems that intelligently derive meaning and context from natural language text. “Natural languages” are languages that are spoken by humans (e.g., English, French, Japanese). Computers cannot, without assistance, distinguish linguistic characteristics of natural language text. Natural language processing systems are employed in a wide range of products, including spell checkers, grammar checkers, machine translation systems, and speech synthesis programs.

Oftentimes, natural languages contain ambiguities that are difficult to resolve using computer automated techniques. Ambiguities come in many forms. Confusable words (e.g. then/than, its/it’s, weather/whether) are one of the biggest sources of grammar errors by users. Possessive/plural types (e.g., kids/kid’s) is another source for ambiguity. A third common example is part-of-speech tagging, such as differentiating whether “produce” is a noun or a verb. A fourth example is word sense disambiguation, such as deciding whether a particular instance of the word “crane” is referring to a bird or a machine.

Many natural language processing problems can be viewed as trying to disambiguate a token into one of a small number of possible labels, based upon the string context in which that token appears. For example, a spell checker may try to decide whether the word “then” or “than” is appropriate in the sentence “I am

1 much smarter *then/than* you are.” A machine translation system may try to
2 determine what the word sense is of the word “line” in the sentence “I am not
3 going to wait in *line*”, so it can more accurately determine what the proper
4 translation is. A speech synthesis program may try to decide whether the word
5 “produce” is a noun or a verb in the sentence “This grocery store has beautiful
6 *produce*”, in order to determine the proper pronunciation for the word.

7 To automatically perform disambiguations, the natural language processing
8 system is provided with linguistic knowledge that it applies to the string context in
9 order to disambiguate. Linguistic knowledge can either be entered manually or
10 learned automatically. Typically, the manual approach has the advantage that
11 people can provide linguistically sophisticated knowledge. Automatic approaches
12 are beneficial in that the linguistic knowledge can be derived empirically from
13 essentially unlimited amounts of data, can be rapidly ported to new domains,
14 languages, or problems, and can be constantly and automatically adapted to a
15 particular individual or subpopulation.

16 To date, automatic approaches have been extremely constrained in the types
17 of linguistic information they can learn. For example, conventional systems
18 automatically learn how to disambiguate words/phrases by learning cues based on
19 whether a specific word appears within a pre-specified window of words from a
20 “disambiguation site” (i.e., the place in the text where the ambiguity to be resolved
21 actually occurs), and what combinations of words and word features (such as part
22 of speech) appear in immediate proximity to the disambiguation site. The
23 contextual words and phrases surrounding the disambiguation site are commonly
24 referred to as the “string context” or simply “string”.
25

1 To provide a concrete example of a token disambiguation problem, a
2 spell/grammar checker may wish to check whether the words “then” and “than”
3 are confused anywhere in a document. Suppose the sentence is:

4
5 I am much bigger *then* you.
6

7 The spell/grammar checker will try to determine whether “then” or “than” is the
8 correct word. It does so by analyzing the string context (e.g., the sentence in
9 which the word appears) and applying its linguistic knowledge to this string
10 context to determine which word is more likely. In this particular example, it may
11 make use of its linguistic knowledge that the word “than” immediately follows a
12 comparative adjective much more often than the word “then”.

13 There are two primary components of a machine learning approach to the
14 problem of token disambiguation based on a string context: (1) the algorithms
15 used for learning and applying the learned knowledge to perform disambiguation
16 and (2) the specification of features the learner is allowed to explore in training.
17 Over the past decade, there have been many different approaches to (1), but very
18 little progress in (2).

19 For confusable word set disambiguation, an article by Golding and
20 Schabes, entitled “Combining Trigram-based and Feature-based Methods for
21 Context-Sensitive Spelling Correction,” *Proceedings of the 34th Annual Meeting of*
22 *the Association for Computational Linguistics*, 1996, describes training a naïve
23 Bayes classifier using as features the set of words that appear within +/- 3 words
24 of the target word and patterns of up to 2 contiguous words and/or part of speech
25 tags around the target word. In an article by Golding and Roth, entitled “A

1 Winnow-Based Approach to Spelling Correction,” *Machine Learning*, Special
2 issue on Machine Learning and Natural Language Processing, Volume 34, pp.
3 107-130, 1999, the authors propose using the Winnow machine-learning algorithm
4 with essentially the same features. In an article by Mangu and Brill, entitled
5 “Automatic Rule Acquisition for Spelling Correction,” *Proc. of the Fourteenth*
6 *International Conference on Machine Learning, ICML'97*, Nashville, Tennessee,
7 1997, the authors describe use of transformation-based learning, again using the
8 same features. In an article by Jones and Martin, entitled, “Contextual Spelling
9 Correction Using Latent Semantic Analysis,” *Proceedings of the Fifth Conference*
10 *on Applied Natural Language Processing*, 1997, the authors propose use of latent
11 semantic analysis as the learning algorithm, and features that include the set of
12 words and contiguous word pairs (bigrams) that appear within a window of +/- 7
13 words of the target word.

14 For word sense disambiguation, an article by Ng, entitled “Exemplar-Based
15 Word Sense Disambiguation: Some Recent Improvements,” *Proceedings of the*
16 *Second Conference on Empirical Methods in Natural Language Processing*, 1997,
17 describes systems for word sense disambiguation that employ two different
18 machine learning algorithms, naïve Bayes and Nearest-Neighbor. In both systems,
19 the features used were: word before, word after, word two before, word two after,
20 the pair of words before, the pair of words after, and the two surrounding words.
21 In an article by Yarowsky, entitled “One sense per collocation,” *In Proceedings of*
22 *the ARPA Human Language Technology Workshop*, 1993, the author proposes
23 using a decision list learning algorithm with a very similar set of features.

24 One attempt at a richer feature set was proposed by Christer Samuelsson,
25 Pasi Tapanainen and Atro Voutilainen in “Inducing Constraint Grammars,”

published in Grammatical Inference: Learning Syntax from Sentences, Lecture Notes in Artificial Intelligence 1147, Springer (L. Miclet and C. del la Huguera eds), 1996. There they propose a system that can learn barrier rules for part of speech tagging. A barrier rule consists of a pair of symbols X and Y and a set of symbols S, and matches a string if that string contains X and Y, with X preceding Y and no symbols from the set S intervening between X and Y.

Despite these efforts, there remains a need for a method for learning much more expressive disambiguation cues. Such a method should be capable of being applied to virtually any problem involving token disambiguation in a string context, and should offer significant performance gains over current state of the art automatic linguistic knowledge acquisition solutions to these problems.

SUMMARY

A linguistic disambiguation system and method creates a knowledge base by training on patterns in strings that contain ambiguity sites. The system is trained on a training set, such as a properly labeled corpus. The string patterns are described by a set of reduced regular expressions (RREs) or very reduced regular expressions (VRREs), which specify features that the training system is allowed to explore in training. The resulting knowledge base utilizes the RREs or VRREs to describe strings in which an ambiguity occurs. In this way, the technique can be applied to virtually any problem involving token disambiguation in a string context.

In the described implementation, the set of reduced regular expressions (RREs) over a finite alphabet Σ is defined as:

- 1 (1) $\forall a \in \Sigma$: “a” is a reduced regular expression and denotes a set $\{a\}$;
2 “a+” is a reduced regular expression and denotes a positive
3 closure of the set $\{a\}$;
4 “a*” is a reduced regular expression and denotes a Kleene
5 closure of the set $\{a\}$;
6 “~a” is a reduced regular expression and denotes a set $\Sigma - a$;
7 “~a+” is a reduced regular expression and denotes the
8 positive closure of the set $\Sigma - a$;
9 “~a*” is a reduced regular expression and denotes the Kleene
10 closure of the set $\Sigma - a$;
11 (2) “.” is a reduced regular expression denoting a set Σ ;
12 (3) “.+” is a reduced regular expression denoting the positive closure of the
13 set Σ ;
14 (4) “.*” is a reduced regular expression denoting the Kleene closure of the
15 set Σ ; and
16 (5) if r and s are reduced regular expressions denoting languages R and S,
17 respectively, then “rs” is a reduced regular expression denoting a set RS.
18

19 It is noted, however, that reduced regular expressions may contain
20 variations or extensions of the above definition.

21 The set of very reduced regular expressions (VRREs) over an alphabet Σ is
22 defined as:
23

- 24 (1) $\forall a \in \Sigma$: “a” is a very reduced regular expression and denotes a set $\{a\}$;
25 (2) “.” is a very reduced regular expression denoting a set Σ ;

1 (3) “.” is a very reduced regular expression denoting a Kleene closure of
2 the set Σ ; and

3 (4) if r and s are very reduced regular expressions denoting languages R and
4 S , respectively, then “ rs ” is a very reduced regular expression denoting a
5 set RS .

6
7 The set of RREs is strictly greater than the set of VRREs. In other words,
8 every VRRE is an RRE but not every RRE is a VRRE.

9 Once trained, the system may then apply the knowledge base to raw input
10 strings that contain ambiguity sites. The system uses the RRE-based knowledge
11 base to disambiguate the sites.

12 13 **BRIEF DESCRIPTION OF THE DRAWINGS**

14 Fig. 1 is a block diagram of an exemplary training computer.

15 Fig. 2 is a flow diagram of a learning process implemented by the training
16 computer of Fig. 1.

17 Fig. 3 is a flow diagram of a transformation sequence learning process that
18 may be implemented as part of the Fig. 2 learning process.

19 Fig. 4 is a flow diagram of a graph construction process that is implemented
20 as part of the Fig. 3 process.

21 Fig. 5 is an example of a portion of a graph constructed from one iteration
22 of the learning process.

23 Fig. 6 is a block diagram of an exemplary computer that implements a
24 trained disambiguation system.

1 Fig. 7 is a flow diagram of a disambiguation process implemented by the
2 Fig. 6 computer to resolve ambiguity sites in input strings.

3 4 **DETAILED DESCRIPTION**

5 This disclosure describes linguistic disambiguation systems and methods
6 that create a knowledge base by training on patterns in strings that contain
7 ambiguity sites and then apply the knowledge base to raw strings to disambiguate
8 such sites. More particularly, the system defines a set of reduced regular
9 expressions (RREs) or very reduced regular expressions (VRREs) that describe
10 particular patterns in the strings. The system creates the knowledge base that
11 utilizes the RREs or VRREs to describe string patterns in which an ambiguity
12 occurs.

13 Once trained, the system may then apply the knowledge base to raw input
14 strings that may contain ambiguity sites. The system uses the knowledge base to
15 disambiguate the sites.

16 Linguistic disambiguation systems may be implemented in many different
17 environments. For example, a system may be implemented in natural language
18 processing systems, which are employed in spell checkers, grammar checkers,
19 machine translation systems, speech synthesis programs, and the like. The system
20 and techniques described herein are not constrained to any particular language or
21 alphabet, but may be applied to any string-based context, such as character strings,
22 strings of symbols, and the like.

23 The following discussion is organized to first define reduced regular
24 expressions and very reduced regular expressions that are employed by the
25 linguistic disambiguation system. The RREs and VRREs specify the types of

1 patterns that the training system is allowed to explore in training. Afterwards,
2 training systems and methods for deriving a knowledge base based on the RREs
3 and VRREs are described, followed by one exemplary implementation of a
4 training system that employs transformation sequence learning. The discussion
5 concludes with a description of an applier system that applies the knowledge base
6 to disambiguate raw strings.

7 8 **Reduced and Very Reduced Regular Expressions**

9 Regular expressions are descriptions of patterns. Regular expressions are
10 conventional and well known to those of skill in the art. Given a string S and a
11 regular expression R, R may or may not match S. A match is defined in one of
12 two ways: (1) either the string S contains the pattern R, or (2) the pattern R
13 completely describes the string S. For discussion purposes, the second definition
14 of match is used herein, unless specified otherwise. It is noted that the
15 disambiguation learner can easily be modified to learn using either definition.

16 Suppose that strings are made by concatenating one or more space-
17 delimited tokens from a set of tokens, Σ . For instance, if the strings are English
18 words, the set of tokens Σ is the set of words in English. A very simple regular
19 expression is “the”. Using the first definition above, this expression would match
20 all strings having the word “the” appear anywhere within.

21 A set of *regular expressions* over a finite alphabet Σ is defined as:

- 22
23 (1) $\forall a \in \Sigma$ (i.e., for all “a” in set Σ), “a” is a regular expression and denotes
24 the set $\{a\}$; and
25

(2) If “r” and “s” are regular expressions denoting the languages R and S, respectively, then (r|s), (rs), and (r*) are regular expressions that denote the sets $R \cup S$, RS and R^* respectively.

There are standard operations that may be used to develop more complex expressions. Standard notations and their corresponding operations are shown in Table 1.

Table 1

<u>Notation</u>	<u>Operation</u>
“+”	One or more. E.g., the regular expression “(the)+” matches all strings with one or more contiguous occurrences of the word “the”.
“*”	Zero or more. E.g., the regular expression “(the)*” matches all strings with zero or more contiguous occurrences of “the”.
“.”	A wildcard that matches any token in alphabet Σ . E.g., the regular expression “the . . . is” matches all strings with the word “the” followed by any three words followed by the word “is”. The regular expression “the .* is” matches all strings with the word “the” followed by zero or more words followed by “is”.
“~”	Not. E.g., the regular expression “~(the) is” matches any string that contains some word other than “the” immediately followed by the word “is”. “the ~(is)* are” would match any string that contains the word “the” followed by zero or more words that are not “is” followed by the word “are”.

1 “|”

Or. E.g., the regular expression “the (big|small) car” matches any string that contains the word “the” followed by either the word “big” or “small”, followed by the word “car”.

2
3
4
5 The above definition is standard for regular expressions. One problem with
6 regular expressions is that they are far too expressive for learning machines to
7 automatically learn. Thus, current machine learning algorithms are not able to
8 feasibly learn regular expression-based knowledge.

9 To overcome this problem, the learning machine for the linguistic
10 disambiguation system described herein employs two less expressive languages,
11 which will be referred to as “reduced regular expressions” (or RRE), and “very
12 reduced regular expressions” (or VRRE). Reduced regular expressions are strictly
13 less powerful than regular expressions, and the very reduced regular expressions
14 are strictly less powerful than reduced regular expressions. The learning method
15 described herein can learn any RRE and VRRE.

16 More particularly, the set of *reduced regular expressions* (RREs) over a
17 finite alphabet Σ is defined as:

- 18
19 (1) $\forall a \in \Sigma$: “a” is a reduced regular expression and denotes a set $\{a\}$;
20 “a+” is a reduced regular expression and denotes a positive
21 closure of the set $\{a\}$;
22 “a*” is a reduced regular expression and denotes a Kleene
23 closure of the set $\{a\}$;
24 “~a” is a reduced regular expression and denotes a set $\Sigma - a$;
25

1 “ $\sim a^+$ ” is a reduced regular expression and denotes the
2 positive closure of the set $\Sigma - a$;

3 “ $\sim a^*$ ” is a reduced regular expression and denotes the Kleene
4 closure of the set $\Sigma - a$;

5 (2) “.” is a reduced regular expression denoting a set Σ ;

6 (3) “ $.^+$ ” is a reduced regular expression denoting the positive closure of the
7 set Σ ;

8 (4) “ $.^*$ ” is a reduced regular expression denoting the Kleene closure of the
9 set Σ ; and

10 (5) if r and s are reduced regular expressions denoting languages R and S ,
11 respectively, then “ rs ” is a reduced regular expression denoting a set RS .

12
13 The set of RREs is a proper subset of the set of regular expressions. It is
14 noted, however, that reduced regular expressions may contain variations or
15 extensions of the above definition. For example, one natural extension is to
16 include classes of symbols. The above definition can be extended to include
17 classes by letting Σ' be the set of symbols in the alphabet and Σ'' be a set of sets
18 of symbols from Σ' , and by letting Σ in the definition above be the union of Σ'
19 and Σ'' . For example, Σ' might be the set $\{a,b,c\}$, and Σ'' might be the set of sets
20 $\{\{a,b\},\{a,c\},\{abc\}\}$. In this case, Σ would be the set $\{a,b,c,\{a,b\},\{a,c\},\{a,b,c\}\}$.
21 An example of an RRE with this set extension is: $a^+ .^* \{a,c\} \{a,b\}^*$, which means
22 a sequence of one or more a 's followed by zero or more characters followed by
23 one symbol from the set $\{a,c\}$ followed by zero or more symbols from the set
24 $\{a,b\}$. The learning algorithm presented below can easily be augmented to handle
25 this inclusion of classes.

The set of *very reduced regular expressions* (VRREs) over an alphabet Σ is defined as:

- (1) $\forall a \in \Sigma$: “a” is a very reduced regular expression and denotes a set $\{a\}$;
- (2) “.” is a very reduced regular expression denoting a set Σ ;
- (3) “.*” is a very reduced regular expression denoting a Kleene closure of the set Σ ; and
- (4) if r and s are very reduced regular expressions denoting languages R and S, respectively, then “rs” is a very reduced regular expression denoting a set RS.

The set of RREs is strictly greater than the set of VRREs. In other words, every VRRE is an RRE but not every RRE is a VRRE. For expository purposes, the following discussion demonstrates how to learn VRREs. The generalization from learning VRREs to learning RREs is straightforward.

The RREs and VRREs describe patterns in strings. The learning machine creates a knowledge base that employs RREs and VRREs to describe virtually any given string in which a disambiguation site occurs. For example, the learning machine can learn a rule for a “then/than” disambiguation site, such as:

Add evidence for the proper word being “then” if the string matches the pattern: *X followed by zero or more tokens followed by a token that is not Y followed by Z followed by one or more tokens that are not Q followed by an R*

where X, Y, Z, Q and R are particular words or features (e.g. parts of speech). Using expression notation from Table 1, this pattern can be expressed as:

$$X.*\sim YZ\sim Q+R$$

Since the learning machine can learn much more expressive concepts than those learned by current state of the art techniques, it can much more precisely acquire the linguistic knowledge necessary to accurately disambiguate tokens based upon properties of the string context in which they appear.

Disambiguation Training System

Fig. 1 shows an exemplary training computer 100 that implements a linguistic disambiguation training system. The training computer 100 has a processing unit 102, volatile memory 104 (e.g., RAM), and non-volatile memory 106 (e.g., ROM, Flash, hard disk, floppy disk, CD-ROM, RAID system, etc.). The training computer 100 also has one or more input devices 108 (e.g., keyboard, mouse, microphone, stylus, etc.) and one or more output devices 110 (e.g., monitor, LCD, speakers, etc.). The computer 100 is representative of many diverse types of computing devices, including personal computers, server computers, and the like. Although generally illustrated as a single computing system, the training computer may be implemented as multiple computers networked together.

The training computer 100 trains on a corpus 120, which is stored in non-volatile memory 106. In the string-based learning context, a corpus 120 can be viewed as an ordered set of strings. The training corpus 120 is assumed to contain strings that are truthfully and accurately presented. The corpus may be a properly labeled training set, or some other source of accurate strings (e.g., a news source).

1 A knowledge base 122 is created as a result of training on the corpus 120.
2 The knowledge base 122 consists of information that utilizes RREs and VRREs to
3 describe the strings in the training corpus, including strings that include
4 ambiguities.

5 The training computer 100 has a disambiguator trainer 130 that is
6 illustrated as executing on the processing unit 102. The disambiguator trainer 130
7 trains on the corpus to automatically derive the information that forms the
8 knowledge base 122. The trainer 130 may implement different learning
9 techniques to produce the information. Examples of possible learning techniques
10 include transformation sequence learning, decision tree learning, decision list
11 learning, learning with inductive logic programming, and genetic algorithm
12 learning. A detailed implementation involving transformation sequence learning
13 is described below in more detail.

14 15 **General Learning Process**

16 Fig. 2 shows a general learning process 200 that is implemented by the
17 training computer 100. The process is implemented in software as computer-
18 executable instructions that, when run on the processing unit 102, direct the
19 training computer 100 to perform the operations illustrated in the blocks.

20 At block 202, the reduced regular expressions (RREs) and/or very reduced
21 regular expressions (VRREs) are defined. The RREs and VRREs essentially
22 specify the types of patterns that the trainer 130 is allowed to explore in training
23 on the corpus. The defining operation may occur during initialization where the
24 training process loads the user-defined RREs/VRREs. In the exemplary
25 implementation, the training system utilizes the RREs/VRREs described above

1 under the heading “Reduced and Very Reduced Regular Expressions” to establish
2 the allowable pattern types for the trainer.

3 At block 204, the training computer 100 reads a training corpus 120. At
4 block 206, the training computer 100 then creates information that uses the
5 RREs/VRREs to properly resolve ambiguity sites in the training corpus based
6 upon the string context in which they appear. As one example of this creation
7 operation, the trainer 130 implements a learning machine that is augmented with
8 the ability to learn with RREs/VRREs. The learning machine may utilize, for
9 example, transformation sequence learning to produce a set of rules based on the
10 RREs/VRREs, or decision tree learning to produce classification trees based on
11 the RREs/VRREs.

12 At block 208, the training computer 100 accumulates and stores the
13 information as a knowledge base 122. Once trained, the knowledge base 122 can
14 be incorporated into products that apply the learned knowledge to resolve
15 ambiguity sites in new strings. Application of the knowledge base to new strings
16 is described below in more detail under the heading “Disambiguation Process”.

17 It is noted that the training computer generates the knowledge base 122
18 automatically. The end result may contain errors or oversimplifications that are
19 very difficult to catch through automated learning models. Accordingly, to cure
20 such errors or oversimplifications, a skilled linguist may opt to read and modify
21 the knowledge base.

22 23 **Exemplary Transformation Sequence Learning**

24 Transformation sequence learning is one example learning technique that
25 may be implemented by the disambiguation trainer 130. The following discussion

1 presents a detailed example of a disambiguation trainer that implements
2 transformation sequence learning.

3 Transformation sequence learning attempts to find the best rule to describe
4 a string in a training set by iteratively applying a sequence of rules, in order, to the
5 string. An exemplary transformation-based learning machine includes:

6
7 (a) A start-state annotator, which assigns an initial label to a string.

8 (b) A sequence of rules of the form:

9
10 *Change the label of a string from m to n if C(string)*

11
12 where C is some predicate over strings.

13 Fig. 3 shows a transformation sequence learning process 300 that may be
14 implemented by the disambiguation trainer 130 in Fig. 1. The process 300
15 provides one detailed implementation of block 206 from Fig. 2. The process 300
16 is preferably implemented in software as computer-executable instructions that,
17 when run on the processing unit 102, direct the training computer 100 to perform
18 the operations illustrated in the blocks.

19 At block 302, the disambiguation trainer 130 removes the labels from one
20 or more strings of a properly labeled training set and applies an initial or “start-
21 state” annotation to each string. The start-state annotation is a preliminary guess
22 as to which label should be applied to a string. The guess may be random, or
23 based on an analysis of the original labels on the string. The start-state annotator
24 applies the annotation based on any computable function over strings. As one
25

1 suitable example, the start-state annotator initially labels all strings with the most
2 frequent label that appears in the training set.

3 At block 304, the trainer 130 identifies the “best” rule to apply to the
4 training set. The process 300 does not depend on a particular goodness measure,
5 but one suitable criterion for determining the “best” rule for each learning iteration
6 will be the rule that results in the greatest improvement in accuracy when applied
7 to the training set.

8 The “best” rule is appended to the end of the learned transformation
9 sequence (block 306) and then applied to the training set (step 308). This learning
10 iteration, represented by blocks 304, 306, and 308, is repeated until a stopping
11 criterion is met, as represented by block 310.

12 To demonstrate the transformation sequence learning process 300, consider
13 the following example of learning reduced regular expressions using a predicate
14 $C(\text{string})$ in the form of: “Does RRE R apply to the string?” More particularly, the
15 learner will learn a sequence of rules of the form:

16
17 *Change the label of a string from l_i to l_j (for $l_i, l_j \in L, L = \{0,1\}$)*
18 *if the string matches reduced regular expression R .*

19 The process is demonstrated in the context of a binary classification case
20 (where $|L| = 2$). It is noted that in binary classification, transformation lists reduce
21 to decision lists, and rules of the form $X \rightarrow Y$ if C (i.e., transform X to Y if C) are
22 equivalent to “Label as Y if C ”. However, the ongoing discussion continues in the
23 transformation-based framework.

24 In this example, the “best” rule is determined in terms of good and bad
25 label changes. The rule “ $l_i \rightarrow l_j$ if C ” will result in a good label change to a string

1 if C applies to that string, the string is currently labeled as l_i , and the proper label
2 for the string is l_j . The rule will result in a bad label change if C applies, the string
3 is currently labeled as l_i , and the proper label for the string is l_i . The goodness of a
4 rule is measured as the number of good label changes minus the number of bad
5 label changes. It is noted, however, that other goodness measures may be
6 employed.

7 To demonstrate the transformation sequence learning process 300, consider
8 the training set in Table 2.

9
10 **Table 2: Initial Training Set**

<u>String</u>	<u>Label</u>	<u>Initial Label Guess</u>
1: a b c	0	1
2: a b b	1	1
3: b a a	1	1

15
16 Since 1 is the most frequent label in the training set, the start-state annotator
17 initially assigns all three training set strings the label 1 (i.e., block 302), as
18 indicated by the initial label guess in Table 2. As a result, string 1 is incorrectly
19 labeled and strings 2 and 3 are correct. After the initial annotation, the trainer
20 iterates through operations 304-310, in each iteration learning a rule whose
21 application will best improve our labeling of the training set.

22 In each learning iteration in process 300, the trainer constructs intermediate
23 results 140 in the form of a graph containing nodes and paths between the nodes.
24 The trainer attempts to find the “best” node in the graph and return edge labels on
25

1 the path from a root node to the “best” node. These edge labels form the learned
2 RRE or VRRE.

3 Fig. 4 illustrates the graph construction process 400. At block 402, the
4 trainer builds a first or root node in the graph, representing the null RRE or VRRE.
5 Using the Table 2 training set of three strings, the root node of the graph contains a
6 corpus position set: $\{(1,0),(2,0),(3,0)\}$. The notation “(x, y)” for a corpus position
7 represents a string number “x”, followed by a position number “y”. Positions are
8 defined as the locations just before and just after a character. For string 1 in Table
9 2, for example, there are four positions 0, 1, 2, and 3, as follows:

10
11
$${}^0{}_a{}^1{}_b{}^2{}_c{}^3$$

12

13 At block 404, the trainer expands the graph from the root node to secondary
14 nodes with each unique path from root node to each secondary node representing
15 an RRE/VRRE and storing at each node the string positions in the training corpus
16 to which that RRE/VRRE can map. A node is expanded by considering all ways
17 the RRE/VRRE represented by that node can be extended by appending a single
18 atomic RRE/VRRE to the end of it, where an atomic RRE/VRRE is any
19 RRE/VRRE that is not formed through concatenation (the concatenation operation
20 is item (5) in the definition of RREs and item (4) in the definition of VRREs).
21 Ideally, the graph is expanded until it is certain that an optimal RRE/VRRE for a
22 particular training iteration is found. However, in practice, the amount or quantity
23 of expansion is an implementation detail that may vary depending upon the corpus
24 size and processing/memory constraints.
25

1 Fig. 5 illustrates a graph 500 that contains root node 502, and paths to
2 several secondary nodes for learning VRREs. The root node 502 contains the
3 corpus position set $\{(1,0),(2,0),(3,0)\}$ and represents the null VRRE. The paths
4 shown from the root node demonstrate ways of expanding the root node by
5 appending atomic VRREs. The paths lead to secondary nodes that contain
6 secondary position sets resulting from the expansions along associated paths.

7 For example, path 504 represents a case in which the letter “b” is
8 considered as the next character after the null VRRE. Only string 3 which begins
9 with a “b” satisfies this path, because strings 1 and 2 each begin with an “a”.
10 Accordingly, path 504 leads to a secondary node 506 that contains the position set
11 $\{(3,1)\}$.

12 Path 508 represents the case in which letter “a” is considered as the next
13 character after the null VRRE. Here, both strings 1 and 2 match this constraint,
14 thus forming a node 508 that contains a position set $\{(1,1),(2,1)\}$.

15 Path 512 represents an “.” (dot-star) operation which can consume any
16 sequence of zero or more tokens. Path 512 leads to a node 514 that contains the
17 position set $\{(1,0),(1,1),(1,2),(1,3),(2,0),(2,1),(2,2),(2,3),(3,0),(3,1),(3,2),(3,3)\}$.

18 Path 516 represents a “.” (or dot) operation, meaning that any string that
19 matches a token in the alphabet permits movement forward one position. This
20 path leads to a node 518 with a position set $\{(1,1),(2,1),(3,1)\}$.

21 Paths 504, 508, 512, and 516 are merely examples. Other paths may be
22 made, such as one that represent “~” operations. In addition, other paths may
23 branch from the non-root nodes to other non-root nodes, as illustrated by paths
24 520, 522, 524, and 526 to nodes 530, 532, 534, and 536, respectively. A path 540
25 from node 534 to a node 550 is also illustrated.

With reference to block 406 in Fig. 4, once the graph is properly grown, the trainer identifies the “best” node in the graph. In this example implementation, for a node state S with corpus position set P and corresponding RRE/VRRE R , the trainer computes the goodness of the rule “ $0 \rightarrow 1$ if R ” as follows:

$$\text{Goodness_0_to_1}(S) = \sum_{(j,k) \in P} \text{Score_0_to_1}((j,k))$$

$$\text{where } \text{Score_0_to_1}((j,k)) = \begin{cases} 1 & \text{if } k = |C[j]| \text{ and } \text{Guess}[j] = 0 \text{ and } \text{Truth}[j] = 1 \\ -1 & \text{if } k = |C[j]| \text{ and } \text{Guess}[j] = 1 \text{ and } \text{Truth}[j] = 0 \\ 0 & \text{otherwise} \end{cases}$$

where $C[j]$ means the j^{th} string of a corpus C . Similarly, the trainer can compute the score for rule “ $1 \rightarrow 0$ if R ”. $\text{Goodness}(S)$ is then defined as:

$$\text{Goodness}(S) = \max(\text{Goodness_0_to_1}(S), \text{Goodness_1_to_0}(S)).$$

In the Fig. 5 graph, both nodes 536 and 550 have a goodness score of 1 for the corresponding $1 \rightarrow 0$ rule. If shorter rules are preferred, the first rule in the transformation list might be “ $1 \rightarrow 0$ if dot* c”.

At block 408, once the best node (or nodes) is identified as having the highest goodness value, the trainer traces back over the path from root node to best node to identify the RRE/VRRE corresponding to that node. Each node in the graph represents an RRE/VRRE, which is derived from the edge labels on the path from root to the node. In Fig. 5, node 536 represents the RRE/VRRE “.* c”,

1 which means all possible strings that match any tokens in the alphabet followed by
2 the character c. Node 550 represents the RRE/VRRE “a . c”, which means any
3 string that begins with an “a” and ends with a “c”.

4 In this case, one of the two top scoring RREs/VRREs, corresponding to
5 nodes 536 and 550, would be added to the transformation sequence and applied to
6 the training corpus (i.e., blocks 306 and 308). Applying the RRE/VRRE to the
7 three strings 1, 2, and 3 in Table 2 changes the guess array with respect to string
8 from a “1” to a “0”, as represented by Table 3.

9
10 **Table 3: Training Set After One Learning Iteration**

<u>String</u>	<u>Label</u>	<u>New Label Guess</u>
1: a b c	0	0
2: a b b	1	1
3: b a a	1	1

15
16 In our example, after applying the rule to the training corpus, all strings A,
17 B, and C are correctly labeled and training terminates. If the stopping criterion
18 were not met, however, the process would apply the learned rule to change the
19 values of the guess array, then create a new graph, find the best state in that graph,
20 and so on.

21
22 Exemplary Pseudo Code For Process 400

23 The following pseudo code implements the graph building process 400 of
24 Fig. 4. To simplify discussion, the code is directed to learning rules with VRREs.
25 The extension to RREs is straightforward. Suppose that, given a training corpus C

1 and for every string $C[j] \in C$, $\text{Truth}[C[j]] \in \{0,1\}$ is the true label of $C[j]$ and
 2 $\text{Guess}[C[j]]$ is the current guess as to the label of $C[j]$. The algorithm for one
 3 iteration of rule learning follows.

4
 5 MAIN PROGRAM:

- 6 (1) Create a root node with corpus position set $S = \{(j,0) \mid j = 1 \dots |C|\}$ and
 7 push this node onto the processing stack (STACK).
 8 (2) While (STACK not empty) {
 9 STATE = pop(STACK);
 10 Push(dotexpand(STATE),STACK);
 11 Push(dotstarexand(STATE),STACK);
 12 $\forall a \in \Sigma$ Push(atomexpand(a,STATE),STACK)
 13 }
 14 (3) Find best state S. Let R be the VRRE obtained by following the edges
 15 from the root to S, outputting each edge label as the edge is traversed.
 16 Return either the rule " $0 \rightarrow 1$ if R" or " $1 \rightarrow 0$ if R" depending on which is
 17 appropriate for state S.

18
 19 dotexpand(STATE) {
 20 create new state STATE'
 21 let P be the corpus position set of STATE
 22 $P' = \{(j,k) \mid (j,k-1) \in P \text{ and } k-1 \neq |\text{Corpus}[j]|\}$
 23 If (P' not empty) {
 24 Make P' the corpus position set of STATE'
 25 Add (STATE,STATE',DOT) to tree edges (in other words,
 26 add an edge from STATE to STATE' labeled with DOT)
 27 return STATE'
 28 }
 29 Else return NULL
 30 }

31
 32 dotstarexand(STATE) {
 33 create new state STATE'
 34 let P be the corpus position set of STATE
 35 $P' = \{(j,k) \mid (j,m) \in P, m \leq k, \text{ and } k \leq |\text{Corpus}[j]|\}$
 36 If ($P' \neq P$) {
 37 Make P' the corpus position set of STATE'
 38 Add (STATE,STATE',DOT*) to tree edges
 39 }

```

1         return STATE'
2     }
3     Else return NULL
4 }
5
6 atomexpand(a,STATE) {
7     create new state STATE'
8     let P be the corpus position set of STATE
9     P' = {(j,k) | (j,k-1) ∈ P, k-1 ≠ |Corpus[j]|, and the k-1st symbol in
10    Corpus[j] is a}
11    If (P' not empty) {
12        Make P' the corpus position set of STATE'
13        Add (STATE,STATE',a) to tree edges
14        return STATE'
15    }
16    Else return NULL
17 }

```

The algorithm above is exponential. There are some optimizations that make it feasible to apply the learning algorithm.

Optimization 1: Prune states that cannot be on the path from root to the best state. Define $\text{GoodPotential_0_to_1}(S)$ as the number of strings s in the training corpus for which $\text{Guess}[s]=0$, $\text{Truth}[s]=1$ and $\exists k : (s,k) \in \text{corpus_position_set}(S)$. $\text{GoodPotential_1_to_0}(S)$ can be similarly defined. Then, $\text{GoodPotential}(S)$ is defined as:

$$\text{GoodPotential}(S) = \max(\text{GoodPotential_0_to_1}(S), \text{GoodPotential_1_to_0}(S)).$$

As the graph is constructed, the trainer keeps track of the largest $\text{Goodness}(S)$ encountered. If that value is X , then for a state S' , if $\text{GoodPotential}(S') \leq X$, it is impossible for any path through S' to reach a state with

1 a better goodness score than the best found thus far. This condition is checked
2 when pushing states onto the stack and popping states off the stack to be
3 processed. If the pruning condition is met, the state is discarded.

4 Optimization 2: Merge states with identical corpus position sets. If a state
5 being pushed onto the stack is already represented by another state with an
6 identical corpus position set, the state need not be pushed onto the stack.
7 Heuristics are employed to decide which of the states with identical corpus
8 position sets we should keep, such as choosing the state with the shortest path to
9 the root.

10 Optimization 3: Prune a state subsumed by another state. If a state S is
11 being pushed onto the stack that already contains a state with a superset of all of
12 the good positions and a subset of all the bad positions of S, then state S need not
13 be processed.

14 Linguistic Disambiguation System

15
16 Fig. 6 shows an exemplary computer 600 that implements a linguistic
17 disambiguation system. The computer 600 has a processing unit 602, volatile
18 memory 604 (e.g., RAM), and non-volatile memory 606 (e.g., ROM, Flash, hard
19 disk, floppy disk, CD-ROM, RAID system, etc.). The computer 600 also has one
20 or more input devices 608 (e.g., keyboard, mouse, microphone, stylus, etc.) and
21 one or more output devices 610 (e.g., monitor, LCD, speakers, etc.). The
22 computer 100 is representative of many diverse types of computing devices,
23 including personal computers, set-top boxes, game consoles, portable computing
24 devices, and the like.
25

1 The computer 600 runs an operating system 620 and an application
2 program 622 that utilizes a natural language processing system. Examples of such
3 application programs include word processors with spell/grammar checkers,
4 speech synthesis programs, and the like. For purposes of illustration, operating
5 system 620 and program 622 are illustrated as discrete blocks stored in the non-
6 volatile memory 606, although it is recognized that such programs and
7 components reside at various times in different storage components of the
8 computer 600 and are executed by the processor 602. Generally, these software
9 components are stored in non-volatile memory 606 and from there, are loaded at
10 least partially into the volatile main memory 604 for execution on the processor
11 602.

12 The knowledge base 122 created by the training computer 100 (Fig. 1) is
13 stored in the non-volatile memory 606. The knowledge base 122 contains
14 information that utilizes RREs/VRREs to resolve ambiguity based upon the
15 strings in which the ambiguity occurs.

16 During execution of the application program 622, the processing unit 602
17 executes a disambiguator 630 that receives a string that contains an ambiguity site.
18 In the context of word processing, the string may have been entered by the user
19 during operation of the program 622. Alternatively, the string may have been read
20 from memory or received from some external source. The disambiguator 630 uses
21 the RREs/VRREs in the knowledge base 122 to resolve the ambiguity site.

22 23 Disambiguation Process

24 Fig. 7 illustrates a disambiguation process 700 that is implemented by the
25 disambiguator 630 of Fig. 6. The process 700 is preferably implemented in

1 software as computer-executable instructions that, when run on the processing unit
2 602, direct the computer 600 to perform the operations illustrated in the blocks.

3 At block 702, the disambiguator 630 receives a string with an ambiguity
4 site. The disambiguator 630 initializes the string if necessary (block 704). In the
5 context of transformation sequence learning, the disambiguator 630 annotates the
6 string with a start-state. The disambiguator then uses the RRE/VRRE knowledge
7 base to disambiguate the ambiguity site in the string (block 706). For
8 transformation sequence learning, the disambiguator applies each RRE/VRRE in
9 knowledge base in sequence to the annotated string. At block 708, the
10 disambiguator outputs the final result.

11 12 **Conclusion**

13 Although the description above uses language that is specific to structural
14 features and/or methodological acts, it is to be understood that the invention
15 defined in the appended claims is not limited to the specific features or acts
16 described. Rather, the specific features and acts are disclosed as exemplary forms
17 of implementing the invention.